

Asymmetric Generalisation of Harmony Triggers

Wendell Kimper
University of Manchester

1 Introduction

In languages with vowel harmony, segments can differ in the extent to which they are able to act as triggers of the process. Kaun (1995) notes that there is a typological asymmetry in segments' ability to trigger rounding harmony — non-high vowels are better triggers than their high counterparts. In Yakut, for example, high vowels are restricted in their ability to trigger rounding harmony, and can only do so when the target is also a high vowel (1a–d), while non-high triggers are not subject to this restriction (1e–h).

(1) Yakut Rounding Harmony (Kaun, 1995)

High Triggers		Non-High Triggers	
a. mur <u>u</u> n- <u>u</u>	'nose-ACC'	e. o <u>o</u> - <u>nu</u>	'child-ACC'
b. tu:n <u>u</u> :g- <u>u</u>	'window-ACC'	f. bo:ro: <u>u</u>	'wolf-ACC'
c. tu:n <u>u</u> :k- <u>ter</u>	'window-PL'	g. oh <u>o</u> : <u>tor</u>	'stoves-PL'
d. tob <u>u</u> k- <u>ka</u>	'knee-DAT'	h. so <u>n</u> - <u>ton</u>	'jacket-ABL'

Kaun also notes that there is a plausible source of phonetic grounding for this asymmetry; non-high vowels are less articulatorily and perceptually extreme along the ROUND and BACK feature dimensions (Terbeek, 1977; Linker, 1982; Kaun, 1995). Lowering the jaw and constricting the lips are somewhat articulatorily antagonistic, and the trapezoidal shape of the oral cavity means that lower vowels are less separated along the FRONT/BACK dimension. Harmony is perceptually advantageous (Suomi, 1983; Kaun, 1995; Gallagher, 2010; Kimper, 2011), and vowels with weaker cues have a greater need for that increase in perceptual salience, so a preference for non-high harmony triggers is consistent with this phonetic motivation.

But what is the relationship between the phonetically-grounded explanation for asymmetries like this and the synchronic grammars which instantiate them? Wilson (2006) proposes instead that learners are systematically *biased* in favour of patterns that are phonetically 'natural'. If this kind of bias forms part of language users' implicit knowledge, it should be accessible in laboratory-learning tasks — Wilson (2006) demonstrates this with velar palatalisation, showing that subjects trained on typologically preferred or dispreferred triggers generalised in ways consistent with the implicational hierarchy found in the typology. However, Moreton & Pater (2012a,b)'s review of the literature on biases in artificial grammar showed robust evidence for *structural* bias (toward simpler patterns), but mixed and inconclusive evidence for *substantive* bias (toward phonetically natural patterns).

In this paper, I use the typological asymmetry found in rounding harmony to investigate substantively grounded bias in learning. If markedness hierarchies form some component of language users' implicit knowledge, evidence that high vowels trigger rounding harmony should also serve as evidence that non-high vowels will also act as triggers; however, evidence that non-high vowels trigger harmony provides no information about the behaviour of high vowels. The prediction, therefore, is that subjects trained on a harmony pattern with high vowels should tend to make *broad* generalisations (extending the pattern to all vowels), while subjects trained on a pattern with mid vowels should tend to make both *both narrow and broad* generalisations.

* Many thanks to Anie Thompson for lending her voice to the stimuli, and to Research Assistants Michaela Hejná, Ziyuan Pan, and Isabel Maccabee for help in running subjects. This work has benefited from insightful discussions with Elliott Moreton, Joe Pater, Andrew Stewart, the University of Manchester Phon-Phon Group, the UMass Amherst Sound Workshop, and audiences at AMP 2015, Yale University, and the University of Utah; all errors and oversights remain my own. This research was funded by a grant from the School of Arts, Languages and Cultures Research Support Fund at the University of Manchester.

2 Methods

2.1 Subjects Data for this experiment come from 67 native speakers of British English, recruited from the University of Manchester and surrounding community; they received either course credit or a £10 Amazon voucher for their time. Subjects were excluded from analysis if they were non-native speakers, reported any speech, hearing, or learning disabilities, or gave null responses on more than 10% of test items — an additional 4 subjects participated but were excluded (all non-native speakers). Subjects were randomly assigned to two experimental groups: 33 subjects were assigned to the *mid-only*, while 34 were assigned to the *high-only* group.

2.2 Stimuli Stimuli consisted of trisyllabic nonce words, composed of a disyllabic stem plus a monosyllabic suffix. The vowel sequences used to compose the stems can be seen in Table 1 — vowels within a stem agreed for both height and colour, but ATR features varied freely. The ATR feature dimension was included to provide a sufficiently rich inventory for feature-based rather than segment-based generalizations (cf. Finley 2008).

	Mid		High		
Front/Unround	ɛ e	e ɛ	ɪ i	i ɪ	ATR-disharmonic
	e e	ɛ ɛ	ɪ ɪ	ɪ ɪ	ATR-harmonic
Back/Round	ɔ o	o ɔ	ʊ u	u ʊ	ATR-disharmonic
	o o	ɔ ɔ	u u	ʊ ʊ	ATR-harmonic

Table 1: Vowel sequences in nonce word stems

Stems were formed by pairing each vowel combination in Table 1 with a randomly-selected non-identical pair of consonants from the set {p, t, k, s, b, d, g, z, m, n}; items which closely resembled existing English words were replaced by re-sampling. A total of 144 stems was created; 14 for each ATR-harmonic vowel sequence and 4 for each ATR-disharmonic sequence. Suffixes were constructed to reflect singular/plural inflection subject to ROUND/BACK harmony; the singular suffix was *-ge/-go* and the plural suffix was *-gi/gu*.

Stimuli were recorded in a sound-attenuated room with a head-mounted microphone, produced by a phonetically-trained female native speaker of North American English. The speaker was instructed to avoid diphthongisation of the +ATR vowels, and to place stress wherever felt most natural (but to use the same stress pattern across all items) — this resulted in primary stress on the initial syllable of the stem, with secondary stress on the final syllable (the suffix). To avoid coarticulatory influence between stem and suffix, they were recorded separately (in a neutral frame sentence). Stems were recorded with a dummy *-gə* suffix, and suffixes were recorded with dummy *CəCə* stems; stems and suffixes were then spliced together (with the boundary occurring during the closure of the suffix [g]).

2.3 Task The experiment used an explicit learning task; subjects were instructed that they would be learning how to form plurals in Martian. Subjects received two rounds of training, with each round consisting of passive listening followed by testing with feedback.

In the first training phase, 48 singular/plural pairs were presented auditorily, accompanied by visual presentation of singular and plural images of various fruits. The *mid-only* group saw items with only mid vowels in the stems, and the *high-only* group saw items with only high vowels in the stems. The singular-plural alternation introduced subjects to suffix vowels of both height classes — the *mid-only* group saw high vowels in the plural suffix, and the *high-only* group saw mid vowels in the singular suffix. This provides evidence to all subjects that both high and mid vowels are included in the inventory of Martian, and that both participate in harmony as undergoers.

In the second training phase, plural forms were presented auditorily with no accompanying image — subjects were prompted to indicate via a button press whether or not the item they had just heard was a correctly-formed plural, and received feedback on the accuracy of their response. All 48 stems from the first training phase appeared in the second; half appeared with the previously encountered (harmonic) suffix, and half appeared with an incorrect (disharmonic) plural suffix; correct/incorrect presentations were

counterbalanced across FRONT/BACK and \pm ATR items. This procedure was repeated for the second round of training, with the same items.

In the test phase, plural forms were presented auditorily (with no accompanying image) and subjects were prompted to indicate via a button press whether or not the item they had just heard was a possible plural form of Martian. Items for the test phase included 24 *old* items (previously encountered stems), 24 *new* items (previously unheard stems of the same height class as training) and 48 *novel* items (items of the stem type withheld during training — high-vowel stems for the mid-only group, and mid-vowel stems for the high-only group). Half of the test items appeared with an incorrect (disharmonic) plural suffix, and correct/incorrect presentations were counterbalanced across HIGH/MID, FRONT/BACK, and \pm ATR items.

The experiment took place in a sound-attenuated room in the University of Manchester Phonetics Lab. Instructions and stimuli were presented using E-Prime 2.0 Professional (Schneider et al., 2002); auditory stimuli were played through circumaural headphones at a comfortable volume, and responses were collected using a Serial Response Box. Subjects were offered brief in-situ breaks between each of the training phases, and a mandatory longer break between the training and test phases. After the experiment concluded, subjects were asked to fill in a short debriefing form.

3 Results

Subjects were classified as *learners* if they performed better than chance on the old items in the test phase, and as *non-learners* otherwise. There was a high failure rate: only 15 subjects in the mid-only group and 12 subjects in the high-only group were classified as learners. There was no statistically detectable difference between the two groups in their performance on old items.

A generalised linear mixed effects model was fitted on proportion correct for responses to *new* and *novel* items only, with random intercepts for item and subject. Fixed effects and their interactions were added one at a time until model fit was no longer improved; sum coding was used throughout, and simple effects were obtained by re-running the model on subsets of the data (uncorrected p-values reported).

Figure 1 shows that, among learners, generalisation to new and novel items exhibited the asymmetry expected under substantive bias: the interaction between item type (new vs. novel) and group (mid-only vs. high-only) was significant ($p < 0.01$), with subjects in the mid-only group extending the harmony pattern significantly to new items at a higher rate than novel items ($p < 0.001$) but subjects in the high-only group showing no such distinction ($p > 0.05$).

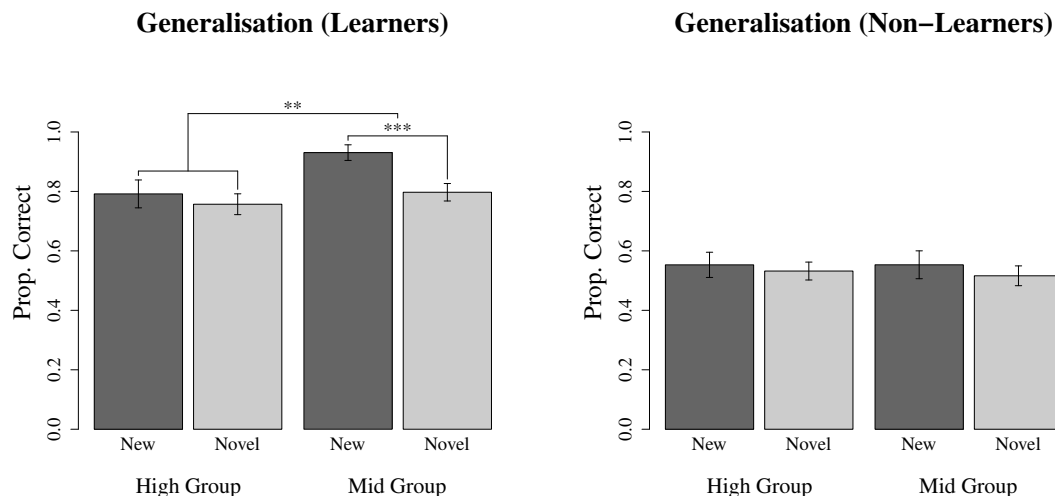


Figure 1: Generalisation to new and novel items, for both learners and non-learners. Error bars represent 95% confidence intervals.

Subjects in the mid-only group showed higher performance overall on the new and novel items than

those in the high-only group ($p < 0.001$). This difference suggests a possible alternative explanation for the asymmetry: could differentiation between new and novel items be a function of general learning success rather than bias?

There was a clear bimodal distribution among learners in performance on the old items: subjects' performance was either below 80% or above 90%, with none in the interval between. On this basis, subjects were classified as either *low performers* (below 80%) or *high performers* (above 90%). If learning success is driving the apparent asymmetry, we should find that high performers in both groups show differences in generalisation between new and novel items, while low performers in both groups should not.

Figure 2 shows that this explanation does not fit the data. While low performers do show symmetrical generalisation for both groups, within the category of high performers we still see the asymmetry predicted under substantive bias: subjects in the mid-only group show greater generalisation to new items than novel items ($p < 0.05$) but the high-only group did not ($p > 0.05$).

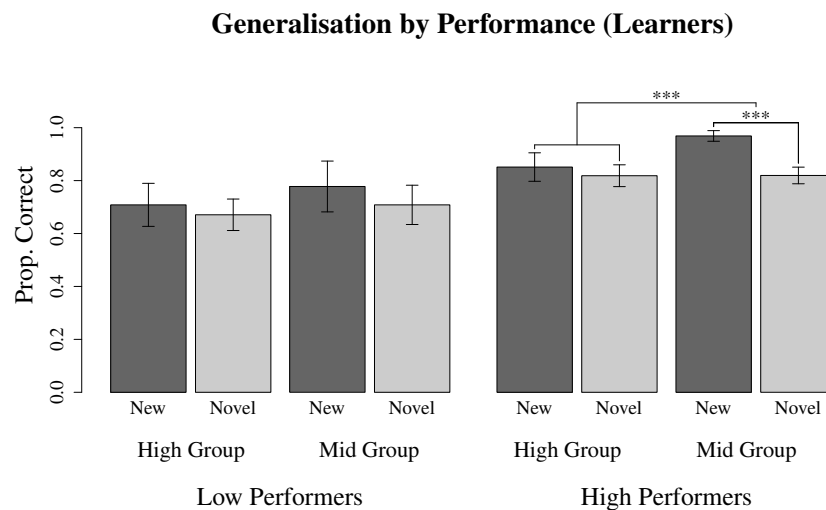


Figure 2: Generalisation to new and novel items for low performers and high performers. Error bars represent 95% confidence intervals.

During debriefing, subjects were asked whether they felt they had identified a pattern and, if so, to describe it. Subjects were classified as successful if the pattern that they described was consistent with the true vowel harmony alternation — Table 2 shows that learners in both experimental groups tended to successfully identify a harmony pattern, while non-learners did not.¹

	Learners		Non-Learners	
	no	yes	no	yes
high-only	3	9	21	1
mid-only	1	14	17	1

Table 2: Successful identification of an explicit harmony rule in debriefing (counts) for both learners and non-learners.

¹ Debriefing also provided some clues to the source of the high failure rate; some non-learners reported that they had been attempting to form rules based on various visual properties of the fruit, and many reported attempting to assign stems to masculine and feminine classes (not entirely surprising, given likely L2 exposure to e.g. French or German during secondary schooling for many subjects).

The rules described by subjects were further classified as *broad* if they mentioned both high and mid vowels, and *narrow* if they mentioned only one height class. Examples of broad and narrow rules are given in (2–3); rules with ambiguous descriptions were classified as ‘unclear’ and set aside.

- (2) **Broad:** *I think the word stems with low vowels - eg. ‘u’ or ‘o’ had the single suffix ‘o’ and the plural suffix ‘u’. Words with ‘i’, ‘a’, or ‘e’ took the ‘i’ plural ending.*
- (3) **Narrow:** *I thought that ‘o’ sounds in the word meant plural was ‘oo’ e.g. nonogoo, & the same with ‘e’ e.g. nenegee.*

Table 3 shows that mid-only learners were roughly split between reporting broad and narrow rules — high-only learners, however, reported *only broad rules*. Figure 3 shows that subjects performance in the test phase of the experiment was consistent with their reported rule types — subjects who reported broad rules (whether high-only or mid-only learners) did not systematically distinguish new and novel items in their performance, but the subjects who reported narrow rules showed a robust difference between new items and novel items.

	broad	narrow	none	unclear
high-only	8	0	3	1
mid-only	5	8	1	1

Table 3: Type of rule identified in debriefing (counts) for learners.

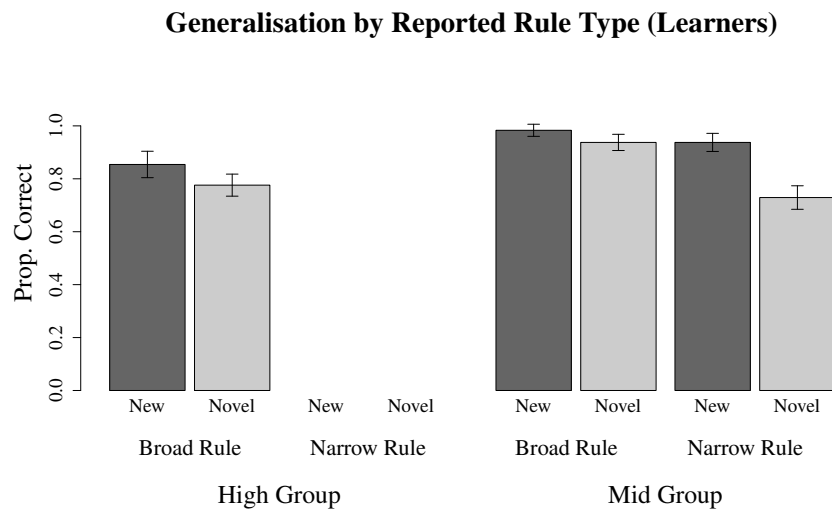


Figure 3: Generalisation to new and novel items for subjects who reported broad vs. narrow rules. Error bars represent 95% confidence intervals.

4 Biased Learning

The results presented in the previous section are consistent with the predictions of substantively biased learning — subjects behaviour mirrors the implicational markedness hierarchy among triggers found in the typology of rounding harmony. This section reports the results of Maximum Entropy (MaxEnt) learning simulations, showing that a learner equipped with substantive bias mirrors the results of human learners more closely than an unbiased alternative.

The learning model used here is not intended to be a direct representation of what the human learners are doing in the lab; there are non-trivial differences between the type of task used in the experiment discussed in

the previous section and the assumptions that underlie probabilistic learning in constraint-based grammars. However, this kind of modelling is still a useful abstraction in this case — demonstrating that a computational implementation of substantively biased learning can be seen as a necessary (but not sufficient) component of an argument that human learners are subject to these biases.

Simulations in this paper used the MaxEnt Grammar Tool (Wilson & George, 2009) to compute constraint weights on the basis of training data; see Goldwater & Johnson (2003) for more on MaxEnt, Hayes & Wilson (2008) for detailed information about the learner, and Wilson (2006) for a demonstration of biased learning using a similar model. Bias in this model is instantiated via regularisation terms μ (which designates a target weight for each constraint) and σ (which specifies a penalty for deviation from that target).

The input data for the learning simulation was a schematised version of the training data used in the experiment above (reducing the items to their vowel sequences) — mid-vowel stems and high-vowel stems were separate, as they were for the human learners. In addition to substantively-grounded bias favouring mid-vowel triggers, the simulations also modelled a variable bias towards breadth in generalisation — the assumption here is that this differs from individual to individual (as evidenced by the split among mid-only learners between broad and narrow generalisations).

Three different learners were tested, each trained on mid-vowel stems and high-vowel stems separately; the constraints and the values for the regularisation terms are given in Table 4. The relevant constraints are above the dashed line at the top of the table; AGREE(V-V) is a general constraint demanding agreement among adjacent vowels, while AGREE(Mid-V) and AGREE(High-V) demand agreement when the trigger (here the final stem vowel) is mid or high respectively. Other constraints are included as distractors.

The target μ for all constraints was set to 0, and biases are implemented by varying σ . Both the *General* and *Conservative* learners implement a substantively-grounded bias favouring mid-vowel triggers — σ is 0.5 for AGREE(Mid-V) and 0.01 for AGREE(High-V). Because smaller values for σ result in greater penalties for deviation from μ (which is 0) and constraint weights must be positive, this should mean that the weight of AGREE(Mid-V) is more easily increased, resulting in a higher likelihood of narrow generalisations based on mid-vowel triggers than on high-vowel triggers. Similarly, the General learner has a σ of 2 for AGREE(V-V), while the Conservative learner's σ is 0.5 — this should mean that the General learner should favour broad generalisations to a greater extent than the Conservative learner. For the *Unbiased* learner, σ for both AGREE(Mid-V) and AGREE(High-V) is set to 0.01, while σ for AGREE(V-V) is set to 1 (a value intermediate between the General and Conservative learners).

Constraint	General		Conservative		Unbiased	
	μ	σ	μ	σ	μ	σ
AGREE (V-V)	0	2	0	0.5	0	1
AGREE (Mid-V)	0	0.5	0	0.5	0	0.01
AGREE (High-V)	0	0.01	0	0.01	0	0.01
AGREE (Front-V)	0	0.01	0	0.01	0	0.01
AGREE (Back-V)	0	0.01	0	0.01	0	0.01
AGREE (ATR-V)	0	0.01	0	0.01	0	0.01
AGREE (RTR-V)	0	0.01	0	0.01	0	0.01
AGREE (i-V)	0	0.01	0	0.01	0	0.01
AGREE (ɪ-V)	0	0.01	0	0.01	0	0.01
AGREE (u-V)	0	0.01	0	0.01	0	0.01
AGREE (ʊ-V)	0	0.01	0	0.01	0	0.01
AGREE (e-V)	0	0.01	0	0.01	0	0.01
AGREE (ɛ-V)	0	0.01	0	0.01	0	0.01
AGREE (o-V)	0	0.01	0	0.01	0	0.01
AGREE (ɔ-V)	0	0.01	0	0.01	0	0.01

Table 4: Constraints and regularisation terms for the three simulated learners.

After training on the schematised input data, both the mid-only and high-only simulated learners were tested on a similarly schematised version of the test items used with the human subjects; using the weights arrived at in learning, the MaxEnt grammar assigns a probability to each item. The *Trained* items were those

vowel combinations that had been part of the input training data (mid-vowel stems for the mid-only simulated learners, and vice versa) and *Novel* items were vowel combinations that the learner had not been trained on (high-vowel stems for the mid-only simulated learners, and vice versa).

Figure 4 shows the probabilities assigned by the Unbiased MaxEnt learner (as before, ‘correct’ responses are those showing vowel harmony). As expected, neither the mid-only nor the high-only simulated learners differentiated between trained and novel items; this is consistent with the human subjects behaviour in the high-only group, but does not mirror human performance in the mid-only group.

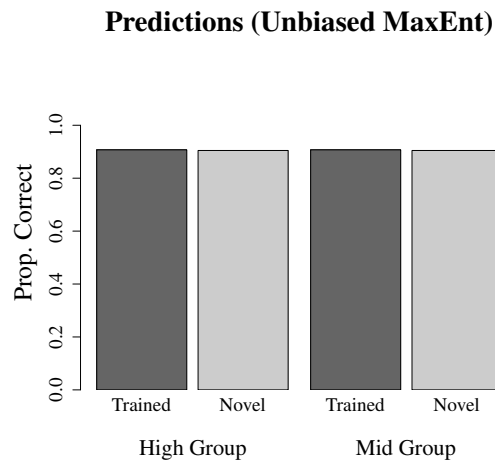


Figure 4: Probabilities assigned to Trained and Novel items by the Unbiased MaxEnt learners.

Figure 5 shows the probabilities assigned by the General and Conservative MaxEnt learners. When trained on high-vowel stems, both the General and Conservative learners formed broad generalisations, making no distinction between trained and novel items. When trained on mid-vowel stems, however, the General learner formed a broad generalisation, making very little distinction between trained and novel items, while the Conservative learner formed a narrow generalisation and did make that distinction.

The assigned probabilities in Figure 5 should be compared with the human subjects’ behaviour as shown in Figure 3 — there is a good qualitative match between the pattern produced by the human subjects and that produced by the simulated learners. The model successfully predicts high-only subjects’ consistent use of broad generalisations, and also successfully predicts the split between mid-vowel subjects between broad and narrow generalisations.

5 Discussion

The results from the experimental task are consistent with substantively-grounded bias in learning, mirroring the implicational markedness hierarchy found in the attested typology, and a simulated learner implementing both substantive bias and a varying generality bias matched the experimental results better than an unbiased learner.

However, some aspects of the task and the subjects’ performance seem to suggest that this result should be somewhat surprising. First, the task used was an explicit learning task — alternations were presented with morphological/lexical context, and subjects were told to identify the pattern governing plural formation in Martian. This kind of explicit task should reflect conscious, domain-general learning strategies rather than the kind of implicit learning that characterises language acquisition (Moreton & Pertsova, 2015). Additionally, van de Vijver & Baer-Henney (2014) found that effects consistent with substantive bias tend to emerge under conditions of uncertainty (e.g. noisy/probabilistic training data). However, the divergence between mid-only and high-only learners was most apparent among those who performed the *highest* on old items in testing.

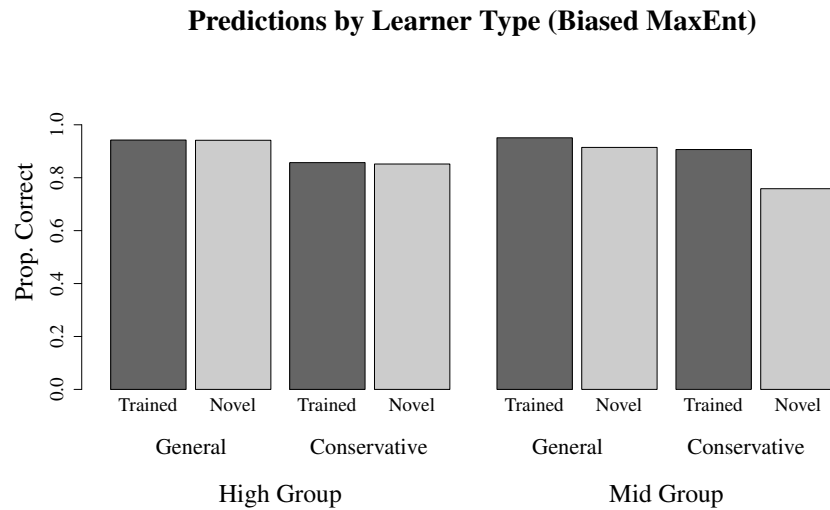


Figure 5: Probabilities assigned to Trained and Novel items by General and Conservative MaxEnt learners.

It's worth taking a moment to consider the task in more detail. Their overt instruction was to learn how to form the plural in Martian; in order to do this, each subject needs to notice that (a) the singular/plural alternation is found in the final vowel (segmentation of the suffix [g] is not crucial), (b) the singular and plural each have more than one surface phonological form, (c) that the form of the affix is predictable from the preceding vowel, and (d) that the basis for prediction is BACK/ROUND agreement. All of these are discernable from the training data, and even the non-learners mostly reported processes or strategies that included reference to both (a) and (b).

However, the Poverty of the Stimulus task leaves a choice between two possible generalisations that are both equally consistent with the explicit instructions and the training data: is the relevant natural class the one containing *all and only* the segments seen in training (narrow) or *any* natural class containing all the segments seen in training (broad)? While the subjects were told that the testing phase would contain some words they hadn't seen before, they had seen no reference to the breadth of their generalisation; it is possible that some subjects did attend to this aspect of the task, but none reported reflecting on this as part of their conscious learning strategy on debriefing forms.²

Embedded in the explicit and overt aspects of the task, the Poverty of the Stimulus paradigm nonetheless includes an implicit component, which in this case requires reference to the phonological categories used in natural language. And it is precisely this component that introduces the uncertainty that van de Vijver & Baer-Henney (2014) found to be most conducive to substantive bias — the training data contains no direct evidence about the breadth of the generalisation.

The support for substantively biased learning in these results comes from the divergent behaviour of the mid-only and high-only subjects. Both groups received training data supporting generalisations of equal complexity, so analytical biases favouring simpler patterns is not an available source of explanation. However, we can see from the performance of the mid-only subjects that there is an additional factor involved in determining the generalisation formed — a tendency towards breadth or narrowness that was modelled separately as an independent bias differing between individuals.

A model of learning with both substantively-grounded bias and a varying bias towards generality predicts, as we can see in MaxEnt learning simulations, that high-only subjects should form broad

² Anecdotal evidence from the performance of students in introductory phonology courses suggests that *even with direct instruction* the choice between broad and narrow natural classes in rule formation frequently evades conscious reflection.

generalisations regardless of their level of generality bias. But this raises the spectre of a possible confound — the possibility that general learners were over-represented in the high-only group cannot be ruled out. Ettliger et al. (2014) found that both declarative and procedural memory can influence the generalisations formed in artificial grammar learning; while there is a general expectation that random assignment to experimental groups should distribute these factors across both groups, future studies should include independent tests of the relevant cognitive dimensions.

6 Conclusion

To summarise, this paper has presented results from an artificial grammar study (in the Poverty of the Stimulus paradigm) which support the claim that learners are biased towards phonetically natural phonological patterns. In the typology of rounding harmony, there is an implicational relationship between high and non-high triggers; non-high vowels, which are comparatively perceptually impoverished along the relevant acoustic dimensions, are preferentially selected as triggers. Subjects in the lab mirrored this implicational hierarchy — those trained on a rounding harmony pattern with high vowels as triggers consistently formed broad generalisations which included mid-vowel triggers, but those trained on mid vowels as triggers were split between forming broad and narrow generalisations. A MaxEnt learning simulation with biases for generality as well as a substantively-grounded bias towards mid-vowel triggers successfully mirrored the behaviour of the human subjects.

References

- Ettliger, Marc, Ann R. Bradlow & Patrick C. M. Wong (2014). Variability in the learning of complex morphophonology. *Applied Psycholinguistics* 35:04, 807–831.
- Finley, Sara (2008). *Formal and cognitive restrictions on vowel harmony*. Ph.D. thesis, Johns Hopkins University.
- Gallagher, Gillian (2010). *The perceptual basis of long-distance laryngeal restrictions*. Ph.D. thesis, MIT.
- Goldwater, Sharon & Mark Johnson (2003). Learning OT constraint rankings using a maximum entropy model. Spenader, Jennifer, Anders Eriksson & Östen Dahl (eds.), *Proceedings of the Workshop on Variation within Optimality Theory*, Stockholm University, 111–120.
- Hayes, Bruce & Colin Wilson (2008). A Maximum Entropy model of phonotactics and phonotactic learning. *Linguistic Inquiry* 39:3, 379–440.
- Kaun, Abigail (1995). *The typology of rounding harmony: An Optimality Theoretic approach*. Ph.D. thesis, University of California, Los Angeles.
- Kimper, Wendell (2011). *Competing Triggers: Transparency and Opacity in Vowel Harmony*. Ph.D. thesis, University of Massachusetts Amherst.
- Linker, Wendy (1982). *Articulatory and acoustic correlates of labial activity in vowels: A cross-linguistic study*. Ph.D. thesis, UCLA.
- Moreton, Elliott & J Pater (2012a). Structure and Substance in Artificial Phonology Learning, Part I: Structure. *Language and Linguistics Compass* 686–701.
- Moreton, Elliott & J Pater (2012b). Structure and Substance in Artificial Phonology Learning, Part II: Substance. *Language and Linguistics Compass* 702–718.
- Moreton, Elliott & Katya Pertsova (2015). Implicit and explicit phonology: What are artificial-language learners really doing? Paper presented at the 23rd Manchester Phonology Meeting.
- Schneider, W., A. Eschman & A. Zuccolotto (2002). *E-Prime User's Guide*. Psychology Software Tools Inc.
- Suomi, Kari (1983). Palatal harmony: A perceptually motivated phenomenon? *Nordic Journal of Linguistics* 6, 1–35.
- Terbeek, Dale (1977). *A cross-language multi-dimensional scaling study of vowel perception*. Ph.D. thesis, UCLA.
- van de Vijver, Ruben & Dinah Baer-Henney (2014). Learners' little helper. Paper presented at the Annual Meeting of the Linguistics Association of Great Britain.
- Wilson, Colin (2006). Learning phonology with substantive bias: An experimental and computational study of velar palatalization. *Cognitive Science* 30, 945–982.
- Wilson, Colin & Ben George (2009). Maxent grammar tool. Software, UCLA. Available for download at <http://www.linguistics.ucla.edu/people/hayes/MaxentGrammarTool/>.

Appendix

Fixed effects:

	Estimate	Std. Error	z value	$Pr(> z)$	
(Intercept)	1.95424	0.21226	9.207	< 2e-16	***
group	-0.92192	0.22710	-4.060	4.92e-05	***
new.novel	-0.73092	0.17679	-4.134	3.56e-05	***
performance	-0.84767	0.22239	-3.812	0.000138	***
atrV2	0.19250	0.07423	2.593	0.009508	**
colour	-0.20479	0.07363	-2.781	0.005415	**
explicit.rule1	0.66988	0.24653	2.717	0.006583	**
explicit.rule2	-0.81856	0.31655	-2.586	0.009713	**
explicit.rule3	0.20819	0.36167	0.576	0.564867	
group:new.novel	0.51557	0.17719	2.910	0.003619	**
group:performance	0.43164	0.20159	2.141	0.032261	*
new.novel:performance	0.43869	0.15673	2.799	0.005126	**
group:new.novel:perf.	-0.40171	0.15673	-2.563	0.010373	*